

Persistent Access Control
or
How to retain control of your information even after distributing it

Paul B. Schneck

MRJ Technology Solutions

(703-277-1618, schneck@mrj.com)

Abstract

A revolutionary invention (patent pending) provides an information owner with the ability to retain control *after distribution* over who may use his information and how it may be used and redistributed. Intelligence reports, imagery, technical data, business information, movies, and audio recordings are but a few examples of the kinds of information that will be protected by this technology. Most importantly, this capability can be delivered in a standard personal computer, allowing users to utilize standard software. This solution, when available as a commercial, off-the-shelf product that provides security and access control is expected to become ubiquitous—a de facto standard for protecting all forms of digital information.

Background

Until recently, protecting some forms of information was relatively simple. Paraphrasing Marshall McLuhan, "The message is in the medium." Copying information required access to the technology used to package the information in the medium. Printed material, photographs, sound recordings, and movies could not be copied easily, could not be copied inexpensively, and certainly could not be copied covertly in an office environment. The widespread availability of xerographic copiers sent a warning signal that technological innovation could remove the barriers to copying. A scant two decades later the ubiquitous personal computer has made it possible for anyone to copy and distribute any digital file. At the same time, the medium is no longer important. Virtually all information types (audio, video, image, text, data, etc.) are now represented as digital data—strings of ones and zeroes—ready for processing by personal computers. The medium on which the bits are delivered, whether tape, magnetic disk, optical disk, telephone line, or other technology, is of no consequence for any subsequent use or processing.

In today's digital world, the source of a string of bits is irrelevant to the personal computer. The string can be processed, copied, or distributed at the touch of a few keys. Even the notion of a "copy" has changed. When dealing with two identical strings of bits (which are intangible) calling one a "copy" and the other the "original" is a distinction

without a difference. That is why, when we purchase or license software we often receive a separate token of ownership: the bits are too easily copied.

Previous methods of control

Before the age of copier technology (xerography, personal computing) several approaches proved reasonably effective in controlling access. In the next few paragraphs we shall take a brief look at these approaches.

“Uncopyable” Media

The earliest examples of art, cave-dwellers’ drawings, are inherently uncopyable. The best we can do is take a picture and reproduce a two-dimensional image, without the texture or the environment. Prior to the invention of the printing press hand printed documents were “effectively” uncopyable. Our use of a hand-written signature to execute a binding contract is a remnant of the view that an individual’s signature is an original that cannot be copied. Even today, so-called “coffee-table books” are testimony to the fact that xerographic copies, although they may capture many aspects, do not capture every aspect of the original. How many times have you seen a coffee-table photocopy? Many personal computer software manufacturers attempted to protect their products by distributing them on “uncopyable” media. As they soon found out, many consumers rose to the challenge and developed techniques for copying these media and, because “bits are bits”, the copies worked just as well as the originals.

Legal and Administrative Controls

The early Hebrew liturgy contains many examples of poetry in which the author’s name is embedded, for example as the first letter of each line. Copying the poem brings the author’s name along. It would be difficult if not impossible to remove the author’s name without damaging or diminishing the original poem. This technique foreshadows the use of “watermarks” or embedded signaling to identify the owner of audio or video information.

The Statute of Anne provided specific printers the sole rights to produce and publish written materials. Modern copyright laws developed from this early example. Today copyright law generally provides that the copyright holder has a property interest in his information. That is, he may act to prevent anyone else from copying or selling the information, subject to the specifics of the law.

The use of administrative techniques is subject to the degree of control that the administrative system can exert on the user. This varies depending on the specifics of each situation. Even in the best cases, administrative controls can be flouted by a criminal or undermined by careless or negligent behavior.

Encryption

The mathematics of encryption allows one to scramble a message so that it can only be unscrambled by another individual possessing the appropriate “key”. This process can be carried out to provide as high a degree of assurance against eavesdropping as is necessary. For example, someone attempting to decrypt a message that was encrypted with a 128-bit key, using one billion personal computers (more than all those ever manufactured) operating at one trillion decryption attempts per second (tens of thousands of times faster than microprocessors that are being designed today) would require over a

trillion years. The universe is about 15 billion years old, only one 60th of the time required!

However, once the message is unscrambled and in the recipient's possession it is just a string of bits. We must trust the recipient to safeguard the information. Encryption can only safeguard the delivery or first access to information. It cannot provide enduring control.

Handling of Classified Information

The security of certain classified information is ensured by the use of an administrative control system and a physical control system. The administrative control system vets all individuals who will have access to classified information. It also puts in place procedures (such as the "two-person rule") for ensuring that a lone individual will not have the opportunity to deviate from the rules. The physical control system provides for secure facilities of processing information and packages and safeguards information in transit between facilities.

The philosophy of the system is that information can only be processed in a secure facility and cannot be removed from the facility without being packaged for secure delivery.

Persistent Access Control

Concept of Operations

Protecting/Packaging

Information that an owner wishes to protect is packaged by the owner and made available to users and potential users. Encryption is one of the packaging steps. This allows the information to be transmitted and distributed without concern that it will be accessed without authorization.

The owner or his representative may send packaged information to users and potential users (the "push" model). Alternatively, the owner may place packaged information on a server in order that users can download it (the "pull" model).

Licenses

Packaged information can be accessed only under the control of a corresponding license. The user or representative prepares a license for each user. Each license specifies the access privileges available to the user. Licenses are themselves protected by encryption (using the public key of the recipient's computer).

Licenses may be distributed with data, to each user on the distribution list or may be prepared and sent in response to users' requests.

Secondary recipients of protected files must contact the owner/representative in order to obtain a license for access

If the owner/representative chooses, a license can authorize specific recipients to distribute copies or derivative products and to provide licenses, perhaps with limited access privileges, to additional recipients.

A small infrastructure is required to validate the authenticity of the public key to be used to encrypt a license. This can be a centralized repository or can be distributed and

implemented as the responsibility of the content owner—who has the greatest interest in ensuring protection. If desired, the computer owner's identity can be associated with a key.

The Computer

A modified computer, containing some additional hardware and software is used to ensure access control in accord with the user's license. The modified computer is compatible with current files. Existing software can be used with current files or with new, protected files.

The modifications include changes to the BIOS to implement an access control mechanism that mediates all input and output operations; a place for storing the computer's private key; and a tamper-detecting enclosure. The tamper-detecting enclosure causes the private key to be erased in the event tampering occurs. When the private key is erased, the computer can no longer access any protected files although it can continue to be used to process current (unprotected) files.

The User

There are no differences in operation when accessing unprotected files. Application programs need not be aware of this new capability. They can continue to operate unchanged when working with unprotected files. When the user's program attempts to access (read) a protected file the system will search for a corresponding license. If none is found, the system will ask the user to locate a license or to terminate the program (similar to what occurs if attempting to read the floppy disk drive when no disk is present). After the license is located the system check the access rights granted to the user. Only if the license allows access will the file be read, decrypted, and made available to the application—which is unaware that these access control operations are taking place. If access is not permitted, the program will receive an error indication or will be terminated. A similar process takes place when the user's program attempts an output (write) operation. The license is consulted to determine whether or not output is allowed and what restrictions may apply (e.g., force "watermarking", restrict to black and white, restrict resolution, etc.). If multiple protected files are in use the system enforces the intersection of the restrictions of all active licenses. This is equivalent to operating at "system high".

"Aware" Programs

The access control mechanism can return a license parameter to a program that makes an inquiry. This allows programs aware of the access mechanism to offer variable levels of features to users, depending on the license that is present.

This new capability allows a software creator to package multiple versions in one release and to license individual capabilities to end-users.

User Identity

As described so far, licenses and access privileges are tied to a specific computer. The implementation can easily allow use of a "smart card" or other token so that licenses can

be tied to a specific user—either on a particular PC or on any of a pool of PCs—allowing intra-enterprise “nomadic computing”.

Conclusion

This invention provides continuing control over access to data, even after they are distributed. Copies of protected files, as well as derivative files (those based on input from protected files) all need authorization of the owner of the original file before a user can obtain access. At no point can a digital copy be made without the owner's authorization.